

Dartmouth College

Dartmouth Digital Commons

Dartmouth College Undergraduate Theses

Theses and Dissertations

6-2-2020

Restoring Humanity to Those Dying Below: An Inquiry Concerning the Ethics of Autonomous Weapons Systems

Juliette A. Pouchol
Dartmouth College

Follow this and additional works at: https://digitalcommons.dartmouth.edu/senior_theses



Part of the [Computer Sciences Commons](#)

Recommended Citation

Pouchol, Juliette A., "Restoring Humanity to Those Dying Below: An Inquiry Concerning the Ethics of Autonomous Weapons Systems" (2020). *Dartmouth College Undergraduate Theses*. 152.
https://digitalcommons.dartmouth.edu/senior_theses/152

This Thesis (Undergraduate) is brought to you for free and open access by the Theses and Dissertations at Dartmouth Digital Commons. It has been accepted for inclusion in Dartmouth College Undergraduate Theses by an authorized administrator of Dartmouth Digital Commons. For more information, please contact dartmouthdigitalcommons@groups.dartmouth.edu.

RESTORING HUMANITY TO THOSE DYING BELOW

An Inquiry Concerning the Ethics of Autonomous Weapons Systems

by

Juliette A. Pouchol

Honors Thesis
Charles Palmer, Advisor
Department of Computer Science
Dartmouth College
June 2nd, 2020

Abstract

Today, autonomous weapons systems promise to make war more precise and effective while removing the human component from the battlefield. With the improvement of deep learning and computer vision, machines will soon be able to navigate and search through contested environments, discriminate between targets, and engage appropriately. The memoirs of drone pilots point to the evolving psychological impact of killing caused by the increase in the amount of empathy and emotional connectedness that drone pilots develop towards their target during the intimate surveillance period. A war fought without “skin-in-the-game” enables drone pilots to become better moral agents and decreases the amount of dehumanization inherent in most types of combat. The decentralized architecture of remote combat produces a more correct following of the Laws of War by reducing the influence of the individual drone pilot. A human on-the-loop system promises to reduce the amount of moral injury incurred by drone pilots and maximize the efficacy and ethicality of the decision-making process at the same time.

Warfare is almost as old as man himself and reaches into the most secret places of the human heart, places where self dissolves rational purpose, where pride reigns, where emotion is paramount, where instinct is king.

– John Keegan, *A History of Warfare*

Table of Contents:

Introduction	5
Chapter One: Autonomous Weapons Systems Towards a More Accurate Implementation of the Laws of War	6
Chapter Two: Morality – Uncodified Governing Principle of War	10
Chapter Three: RPAS - Overcoming Our Biological Inhibitions Against Killing	13
Conclusion	17

Introduction

The invention of the Gatling gun marked a dramatic shift in the landscape of war. This predecessor of the machine gun increased the deadliness of the individual soldier, allowing him to fire munitions of high caliber faster and farther than was previously possible. Richard Gatling, the gun's inventor, wanted to preserve lives by reducing the number of soldiers fighting on the battlefield, but he could not have anticipated the increase in the deadliness of war that ensued. Although the weapon did serve its purpose of reducing the number of soldiers needed on the Gatling-equipped side, the death toll skyrocketed on the rivals' end. Today, fully autonomous weapons systems promise to remove human presence from battlefields altogether – to “*select and engage* without further intervention from a human operator.”¹ With the human entirely out-of-the-loop, *fully* autonomous weapons systems will observe, decide, and act without any human input, free to decide the strategy for achieving a given goal. Although the state of current technology does not yet empower the implementation of fully autonomous weapons systems, this impending reality has raised many questions and concerns. Semi-autonomous weapons system like the Predator Drone, have some degree of autonomy, but retain the human-in-the-loop component, which means that a human is still making the decision to act. For the remainder of this thesis, I will distinguish between autonomous weapons systems (AWS, or semi-autonomous weapons systems) and the purely hypothetical *fully* autonomous weapons systems. AWS promise a more precise and effective war by deploying the individual soldier from behind the safety of her computer monitor. With the increase in AWS usage, the debate surrounding what precisely counts as a “just war” has produced incredibly divergent opinions.

Led by Berkeley Professor Stuart Russel, a campaign to ban the use of AWS conjectures that, similarly to what happened with the Gatling gun, AWS could make war more lethal². Drone strikes have become increasingly ubiquitous, raising concerns about the morality of this new method of waging of war. One objection to taking the human out of the loop claims that in order to handle morally complex situations, it is essential that the agent contextualize the ethical problem. In the context of war, this means recognizing an opponent's humanity and seeing your own desires, emotions, and fears, reflected back at you. In other words, this means having skin-in-the-game. Proponents of this thesis have argued that the Laws of War³ cannot alone guaranty a just war, because moral rules are not codifiable. According to this view, skin-in-the-game is essential to leading a just war because it allows human soldiers to violate the Laws of War and oppose authority if either would cause them to overstep their own moral line. This can mean anything from deciding to spare a child combatant⁴, to refuting the leadership of a genocidal dictator. Russell and others contend that because a machine follows a pre-determined set of rules, “legal” and “moral” are synonymous to it.

¹ Stuart Casey-Maslen et al., “Drones and Other Unmanned Weapons Systems under International Law,” (Leiden: Brill Nijhoff, 2018), 123.

² See, “Open Letter on Autonomous Weapons,” Future of Life Institute, accessed April 29, 2020, <https://futureoflife.org/open-letter-autonomous-weapons/?cn-reloaded=1>.

³ See, *Department of Defense Law of War Manual* (Washington, DC: General Counsel of the Department of Defense, 2015)).

⁴ Paul Scharre, “Army of None: Autonomous Weapons and the Future of War” (New York: W.W. Norton & Company, 2019), 2-4.

Unlike research that has been done in the past, my thesis will evaluate AWS as an alternative to the dehumanization consistently found in a war fought with skin-in-the-game. The first chapter of this thesis will briefly present the three principles of the Laws of War: *necessity*, *proportionality*, and *distinction*, and demonstrate the ways in which AWS might comply, or fail to comply, with those principles. I will then briefly expose the anti-codifiability thesis, which states that because morality is contextual and inexhaustive, it does not translate into programming form. Rather than attempt to refute this argument, I will show that morality, as defined in this manner, is not easily or commonly exercised by the traditional soldier. I will make my claim by evaluating the role of soldiers' biological predispositions against killing in preventing a just war from being fought. Finally, Paul Scharre's and other drone pilots' memoirs will provide the basis for an analysis of the effects of removing a combatant's "fear of death." These memoirs contain evidence that delegating certain high-stakes and high-pressure decisions to machines could make war more precise, effective, and ultimately result in less incidental bloodshed. AWS are not limited by human soldier's physiological and psychological constraints: they come in all shapes and sizes, can loiter for days performing reconnaissance, and execute risky operations. Ultimately, I conclude that AWS are an opportunity to surpass our own shortcomings in war, by removing the human component from the battlefield. For the purpose of this thesis, I will remain within the realm of available evidence; extreme scenarios such as Stuart Russell's "slaughter bots" will be out of the scope of this investigation. I will use the memoirs of drone pilots and other scholarly works to investigate the morality of a remote war, and, having made my claim that the human component should be removed from the physical battlefield, I will attempt to home in on the optimal role for the human component within the war.

Chapter One: Autonomous Weapons Systems Towards a More Accurate Implementation of the Laws of War

Given present-day technological advancements, AWS will soon be able to abide by the Laws of War with un-precedented precision and speed. The Laws of War⁵ (or *Jus in Bello*) are a body of military governance that regulates army conduct during a conflict. This globally endorsed set of rules provides the perfect framework on which to build a goal-driven AWS because it is universally recognized and provides a legal description of a just war. Principally, the Laws of War demand that all parties follow three guiding principles: *necessity*, *proportionality*, and *distinction*. Necessity demands that all potential targets have military value. Distinction requires that parties distinguish between legitimate targets and civilians. Finally, if the first two precepts have been fulfilled, the principle of proportionality demands that the anticipated collateral damage be proportional to the military value of the target. The Laws of War prohibit weapons that cannot be used fully in accordance with these principles, including indiscriminate weapons and weapons which are of a nature to cause "superfluous injury of unnecessary suffering."⁶ As a result of this,

⁵ See, *Department of Defense Law of War Manual* (Washington, DC: General Counsel of the Department of Defense, 2015)).

⁶ Stuart Casey-Maslen et al., "Drones and Other Unmanned Weapons Systems under International Law," (Leiden: Brill Nijhoff, 2018), 100.

many “unlawful” weapons have been banned, including, but not limited to, land mines, poisonous gases, blinding laser weapons, and biological weapons.⁷ Many researchers believe that AWS belong in this category; this chapter seeks to expose the present day uses of and advancements in AWS. In a 2013 interview, Peter Maurer, the president of the International Committee of the Red Cross, stated that “from the perspective of international humanitarian law, any weapon that makes it possible to carry out more precise attacks and helps avoid or minimize incidental loss of civilian life, injury to civilians, or damage to civilian objects, should be given preference over weapons that do not.”⁸ Given this, my aim is not to show that AWS can *perfectly* follow the Laws of War, but rather that they can implement them with a consistently higher degree of accuracy than human soldiers and manned weapons.

The principle of distinction demands that “parties to an armed conflict distinguish between civilian persons and civilian objects on the one hand, and combatants and military objectives on the other, demanding they direct their operations only against combatants and military objectives.”⁹ Conforming to this principle depends primarily on an agent’s information-gathering and analyzing effectiveness. Although AWS capabilities do not yet extend to the differentiation of individual persons, computers have been able to identify cooperative targets for quite some time¹⁰. Cooperative targets emit a secure signal and can be contrasted with non-cooperative targets, which do not broadcast their location, and are therefore harder to track. For example, the Israeli Harpy drone¹¹ can loiter over a predefined area for hours on end, searching for enemy radar emissions. Search-and-attack loitering munitions like Harpy or Raytheon’s JSM/NSM cruise missile¹² have a pre-programmed target which they autonomously find and identify before engaging it, all performed without human intervention. Identifying a non-cooperative target is harder to do than homing in on a radar signal, and therefore demands more advanced technological capabilities.

Paul Scharre, Senior fellow and Director of the Technology and National Security Program at the Center for a New American Security, identifies three requirements for building an autonomous-targeting weapon – “the ability to maneuver intelligently through the environment to search; the ability to discriminate among potential targets to identify the correct ones; and the ability to engage targets, presumably through force.”¹³ Scharre’s third requirement is already implemented today – remotely piloted aircrafts (RPAs) like the Reaper Drone or the Predator Drone are a perfect example of this, with their ability to effectively engage and destroy targets¹⁴. As for the ability to navigate and search surroundings, deep learning neural networks have significantly improved computer vision over the past few years. In 2015 Microsoft’s deep neural network technology performed object-recognition tasks with an error rate of 3.57 percent, beating

⁷ Stuart Casey-Maslen et al., “Drones and Other Unmanned Weapons Systems under International Law,” (Leiden: Brill Nijhoff, 2018), 100-101.

⁸ “The Use of Armed Drones Must Comply with Laws,” ICRC, May 10, 2013, <https://www.icrc.org/en/doc/resources/documents/interview/2013/05-10-drone-weapons-ihl.htm>.

⁹ Stuart Casey-Maslen et al., “Drones and Other Unmanned Weapons Systems under International Law,” (Leiden: Brill Nijhoff, 2018), 101.

¹⁰ Paul Scharre, “Army of None: Autonomous Weapons and the Future of War” (New York: W.W. Norton & Company, 2019), 84-88.

¹¹ See, [iai.co.il](https://www.iai.co.il/p/harpy), accessed May 27, 2020, <https://www.iai.co.il/p/harpy>.

¹² See, Slijper, Frank, and Pax (Organization), “Slippery Slope: The Arms Industry and Increasingly Autonomous Weapons”, (2019. <https://www.paxforpeace.nl/publications/all-publications/slippery-slope>), 17-18.

¹³ Paul Scharre, “Army of None: Autonomous Weapons and the Future of War” (New York: W.W. Norton & Company, 2019), 123.

¹⁴ Bergen et al. “Drone Wars: Transforming Conflict, Law, and Policy” (New York, NY: Cambridge University Press, 2015), 119.

humans at this task by over 1.5%¹⁵. Visual-aided navigation (tracking of the “optical flow”¹⁶) paired with inertial measurement units (sensing “changes in velocities”¹⁷) have enabled drones to explore the indoors and outdoors without the help of GPS. Using the positions and velocities of the objects around it, drones can form an internal representation of their surroundings. Shield AI’s Nova class of artificially intelligent unmanned aerial systems (AI-UAS) are an example of fully autonomous indoor and outdoor navigation that operates even in cluttered, pitch dark, GPS- and comms-denied environments¹⁸. Developing AWS’ visual capabilities is necessary for distinguishing between a legitimate military target and a civilian target, which requires identifying attributes like “clothing, activity, age, and gender” for a person, and “movement, size, shape, color, spacing, and speed”¹⁹ for an object.

Scharre’s second capability, target recognition, is by far the most challenging to develop in AWS, because the Laws of War provide a very broad definition of what constitutes a legitimate target. Indeed, they define the latter as a combatant who is currently participating in the aggression, or an object of military value²⁰. However, the archetype of the legitimate combatant is inconstant and contextual, so it is nigh-impossible to program a perfectly accurate representation into an AWS, just as it is extremely difficult for soldiers to navigate the fog of war. Although the Geneva Conventions require that legitimate combatants bear a “fixed distinctive insignia recognizable at a distance in order to distinguish themselves from non-combatants,”²¹ this rule is not followed in most modern combats, which makes it difficult to know whether a person is a civilian or a legitimate combatant at that moment. In addition, all enemies do not deserve the same treatment: for example, combatants can become “hors de combat” (if they are wounded, or surrender) at any time and should not be targeted. Current technologies for recognizing non-cooperative targets include synthetic aperture radars (SAR)²² that generate grainy pictures from overhead radar pulses. SAR produces images that are very hard to analyze for present-day advanced target recognition (ATR) algorithms and has mainly been used for finding and identifying cooperative targets²³. The Defense Advanced Research Projects Agency (DARPA) is currently working on a program called “Target Recognition and Adaptation in Contested Environments” (TRACE) to improve current ATR algorithms by mobilizing the progress made in computer vision using deep neural networks²⁴. From a hardware perspective, drones present an opportunity to better adhere to the principle of distinction, with their ability to follow targets for long periods of time collecting tactical information. An RPA’s size, speed, and flight-capabilities make it an ideal candidate for navigating the ‘fog of war’ present in all kinds of warfare.

¹⁵ Paul Scharre, “Army of None: Autonomous Weapons and the Future of War” (New York: W.W. Norton & Company, 2019), 87.

¹⁶ Bergen et al. “Drone Wars: Transforming Conflict, Law, and Policy” (New York, NY: Cambridge University Press, 2015), 122.

¹⁷ Ibid., 123.

¹⁸ See, “Nova,” Shield AI, accessed May 4, 2020, <https://www.shield.ai/nova> for an action video of Nova AI-UAS.

¹⁹ Bergen et al. “Drone Wars: Transforming Conflict, Law, and Policy” (New York, NY: Cambridge University Press, 2015), 103.

²⁰ Stuart Casey-Maslen et al., “Drones and Other Unmanned Weapons Systems under International Law,” (Leiden: Brill Nijhoff, 2018), 102-104.

²¹ Bergen et al. “Drone Wars: Transforming Conflict, Law, and Policy” (New York, NY: Cambridge University Press, 2015), 124.

²² Paul Scharre, “Army of None: Autonomous Weapons and the Future of War” (New York: W.W. Norton & Company, 2019), 86.

²³ Ibid., 86.

²⁴ Ibid., 84-88.

In their 2009 work²⁵, professors and roboethicists Ronald Arkin, Patrick Ulan, and Brittany Duncan proposed a model for building an ethical AWS. At the core of their model is the Ethical Governor, which calculates the lawfulness of the lethal actions generated by the AWS. The set of constraints referred to by the governor is derived from the Laws of War (stored in long-term memory) and the rules of engagement specific to a particular mission (stored in short-term memory). The evidential reasoning component of the system takes in the perceptual data gathered by the system and converts it into logical affirmations so that these can be used to query the constraints database. These logical affirmations describe (a) the target, (b) the surroundings and (c) the chosen behavioral response. In order to judge whether or not an action is permissible, the governor retrieves the set of constraints from short- and long-term memory and divides them into two categories: $C_{\text{forbidden}}$ and C_{obligate} . Constraints are derived from the principles of distinction, necessity, proportionality, and others. The interpreter must then “evaluate all of the constraints in $C_{\text{forbidden}}$ to be *false*”²⁶; evaluating even one $C_{\text{forbidden}}$ constraint to be true means that the lethal behavior is not permissible. If this condition is satisfied, at least one constraint in C_{obligate} must evaluate to *true* in order to proceed (i.e. the lethal behavior must be necessary). If either constraint ($C_{\text{forbidden}}$ or C_{obligate}) is not satisfied, then lethal action is restrained, and the AWS is instructed to continue its mission without engaging the target. If both constraints are satisfied, the governor must now ensure that the expected damage is proportional to the expected military value (principle of proportionality). The collateral damage estimator does this by “optimizing the likelihood of target neutralization while minimizing any potential collateral damage that would result from engaging the target with lethal force.”²⁷ This process is performed by the collateral damage estimator, which attempts to minimize damage by performing a risk analysis on every possible weapon at every position that does not violate $C_{\text{forbidden}}$ and choosing the one that produces the least amount of damage. Collateral damage is calculated based on expected infrastructure damage and number of casualties. We can imagine modifying this algorithm to take other variables into account, for example damage to the environment. If the minimum damage is too high given the military necessity level of the target, the lethal action is aborted, else, the AWS has permission to engage its target.

What this chapter shows is that current research and advancements in AWS technology are trying to produce a more just war by enabling a more precise following of the Laws of War. There is no scenario in which a human soldier with skin-in-the-game could perform the calculations achieved by Arkin et al.’s ethical governor within the time constraint of war. Although their architecture is purely experimental, there is no reason why we could not soon see a real implementation of a similar model. Such an enhancement of the process of identifying, validating, and engaging potential targets would dramatically reduce the casualties of war. If AWS do indeed have the potential to reduce the costs of war, then, according to the International Humanitarian Laws, they should be adopted. The following chapter identifies potential philosophical reasons why an AWS could not wage a just war entirely on its own. In particular, it presents the argument that a just war necessitates that human soldiers have skin-in-the-game.

²⁵ See, Ronald C. Arkin et al, “An Ethical Governor for Constraining Lethal Action in an Autonomous System:,” (Fort Belvoir, VA: Defense Technical Information Center, January 1, 2009).

²⁶ Ronald C. Arkin et al, “An Ethical Governor for Constraining Lethal Action in an Autonomous System:,” (Fort Belvoir, VA: Defense Technical Information Center, January 1, 2009), 4.

²⁷ *Ibid.*, 5.

Chapter Two: Morality – Uncodified Governing Principle of War

Many thinkers believe that if morality and empathy are inherently *human* qualities, armed robots should not be deployed onto the battlefield, even if they are fully equipped to adhere to the rules of war. The impetus for this argument is that *legality* and *morality* are essentially different concepts. In their recent work²⁸, Duncan Purves, Ryan Jenkins, and Bradley Strawser define three categories of mistakes that a machine can make. The *empirical* mistake comes from a failure to adequately identify the facts of a situation. Machines make *moral* mistakes when they are in possession of all the facts, but they are unable to come to the right normative decision. Finally, the system can make a *practical* mistake by reacting inappropriately given an apt normative decision²⁹. The improvement of computer vision and weapon precision will likely enable machines to make fewer *empirical* and *practical* mistakes than human soldiers. However, Purves et al. argue that machines will always make more *moral* mistakes than humans. The reason for this is what they call the anti-codifiability thesis: “the true moral theory could [not] be captured in universal rules that the morally uneducated person could competently apply to any situation.”³⁰ This view is based on notions of Aristotelian rationality, which holds that a morally “right” decision is contextual, interpretative, and dependent on the particulars (observed from a subjective point of view) of a dilemma or situation. If this is true, the ability to be morally rational is an inherently human quality that cannot be carried out by rule-based logic³¹. Given this, Purves et al. believe that as long as AI is the “product of a discrete list of instructions provided by humans,”³² machines will never possess the kind of moral judgement required for moral agency.

Philosophers talk about skin-in-the-game as an essential component for leading a just war because humans have a strong biological inhibition against killing that prevents us from readily taking a fellow human being’s life. The concern is then that machines, which may not feel empathy towards humans, will not be restrained by any reticence to harm. In his 2014 book, Lieutenant Colonel Dave Grossman writes that “looking another human being in the eye, making an independent decision to kill him, and watching as he dies due to your action combine to form one of the most basic, important, primal, and potentially traumatic occurrences of war.”³³ According to Grossman, even fighter pilots, who are in the air and therefore further away from their targets, hesitate to pull the trigger: “when it came time to kill, they looked into the cockpit at another man, a pilot, a flier, one of the ‘brotherhood of the air’, a man frighteningly like themselves; and when faced with such a man it is possible that the vast majority simply could not kill him.”³⁴ Overcoming this psychological predisposition against killing requires intense pressure from authority and can

²⁸ See, Duncan Purves, Ryan Jenkins, and Bradley Strawser, “Autonomous Machines, Moral Judgment, and Acting for the Right Reasons,” (Ethical Theory and Moral Practice 18, no. 4 (August 2015)).

²⁹ Ibid., 859.

³⁰ See, Duncan Purves, Ryan Jenkins, and Bradley Strawser, “Autonomous Machines, Moral Judgment, and Acting for the Right Reasons,” (Ethical Theory and Moral Practice 18, no. 4 (August 2015)), 855-856.

³¹ Ryan C. Jenkins et al., « Who Should Die?: the Ethics of Killing in War”, (New York, NY: Oxford University Press, 2018)., 169.

³² Duncan Purves, Ryan Jenkins, and Bradley Strawser, “Autonomous Machines, Moral Judgment, and Acting for the Right Reasons,” (Ethical Theory and Moral Practice 18, no. 4 (August 2015)), 857.

³³ Dave Grossman, “On Killing: the Psychological Cost of Learning to Kill in War and Society,” (New York, NY: Open Road Integrated Media, 2014), 55.

³⁴ Ibid., 54.

lead to soldiers suffering from debilitating moral injury, having crossed their own moral boundary by doing something that violated what they considered to be right. In another chapter of his book, Grossman talks about the “decrease in reality”³⁵ that accompanies distancing oneself from an assault. When soldiers are operating from a removed location, it becomes harder to empathize with their target, which lowers emotional engagement and psychological inhibitions against lethal action. Armin Krishnan calls “adiaphorization” the process of disassociating means from ends when soldiers lose sight of the consequences of their actions. If Grossman and others are right that resistance to killing is directly proportional to the physical distance from the target, then drone pilots should experience very little difficulty to kill.

Distance from the target, however, is not the only way of overcoming one’s biological inhibitions to kill. In his book *Less Than Human*, Professor David Livingstone Smith investigates killology on the battlefield and explains the causes and effects of dehumanization in war. He defines dehumanization as “the act of conceiving of people as subhuman creatures rather than as human beings.”³⁶ When soldiers dehumanize their opponents in war, they are not just making a comparison between them and something less than human, they actually *believe* that their enemy is sub-human and should be treated accordingly. This results from the notion of their being a “great chain of Being”³⁷ with God at the top and lifeless things at the bottom. Although this theory is completely unscientific and disproved by Darwinism, it is still present in today’s rhetoric, as we talk about certain organism being higher or lower than others on an imaginary scale, generally with humans at the top of it. Psychoanalyst Erik Erikson talks about *cultural pseudospeciation*³⁸ to denote the formation of distinct identities and cultures within human race, each with a tendency to believe they are superior from the rest. Pseudospeciation paired with the concept of a “great chain of Being” creates an imagined hierarchy within the human race and is essential to the notion of dehumanization because it results in the lowering of certain groups to a sub-human status. Being able to conceive of the enemy as “sub-human” acts as a “psychological lubricant, dissolving [soldiers’] inhibitions and inflaming [their] destructive passions, [...] empower[ing] [them] to perform acts that would, under other circumstances, be unthinkable.”³⁹ In high pressure situations, fear for one’s life, coupled with military propaganda and intense conditioning, can turn into hate and dehumanization of the enemy.

When philosophers and scientists talk about machines “dehumanizing” war, they are defining the term in a very different way than Smith does. Reducing someone or a group of people to a set of numbers and rules (i.e. lines of code in a program) takes away that person’s or group’s *individuality*, but not necessarily their *humanity*. The example most commonly used to illustrate the difference between these two states of mind is the doctor in her operating room – in order to do her job competently she needs to regard the patient on her table as mere flesh and bones. In doing so, the surgeon disregards her patient’s individuality but does not consider her patient as sub-human. Although the process of dehumanization is not fully understood, this form of self-deception enables soldiers to believe that they are morally justified in committing atrocious acts against their peers. This phenomenon is further reinforced by group solidarity and camaraderie found in military units – soldiers have often adopted dehumanizing expressions to designate the

³⁵ Dave Grossman, “On Killing: the Psychological Cost of Learning to Kill in War and Society,” (New York, NY: Open Road Integrated Media, 2014), 109.

³⁶ *Ibid.*, 26.

³⁷ *Ibid.*, 39.

³⁸ *Ibid.*, 63.

³⁹ *Ibid.*, 13.

enemy. We see instances of this in all kinds of conflicts throughout history. During the holocaust Jews were thought of as “rats” and “cockroaches” by German soldiers and experimented on in horrific ways like the lowliest of lab rats. The Allied forces during WWII called the Nazis “dogs,” and the Japanese were equated to “monkeys, apes, or rodents.”⁴⁰ The Japanese dehumanized their Chinese rivals, who they considered “chancorro,”⁴¹ meaning sub-human or bug-like. Their foreign enemies were often represented with devilish features, like horns, and they were called “monsters” and “devils.” Dehumanization is not a thing of the past; some present-day examples include US soldiers posing with dead Afghan soldiers’ body parts as if they were game animals, or drone pilots referring to their targets as “ants” on a screen⁴².

The anti-codifiability thesis assumes that our capacity to rationalize remains intact in situations of high pressure, like during a war. If this were the case, then Aristotelian rationality would be valuable in order to properly weigh the life of another human being. Skin-in-the-game plays into this thesis because, ideally, it immerses the soldier fully into the scene and provides her with the necessary subjectivity and psychological connectedness to make a correct moral judgement. This skill is essential to preserving humanity when taking a fellow soldier’s life. These two ideas work together to create the ideal soldier, better equipped than machines to make moral judgements. In a recent podcast, Paul Scharre acknowledges a gap between “the reality of war and the fantasy that society tells itself to allow us to send men and women off to war.”⁴³ He highlights the life-altering PTSD that many army veterans suffer from, frequently resulting in divorce, alcohol abuse, and sometimes suicides. In a 2006 report⁴⁴, US Soldiers deployed in Iraq demonstrated that they had at best a very feeble grasp of how to behave ethically and lawfully in war. Only 47% of Soldiers agreed that “all non-combatants should be treated with dignity and respect.”⁴⁵ They were also asked whether or not they insulted non-combatants (yes - Soldiers 28%), and damaged or destroyed civilian objects unnecessarily (yes - Soldiers 9%)⁴⁶. Moreover, around 40% of Soldiers who reported insulting non-combatants also said that in the moment they were experiencing high levels of anger and had also known traumatic combat experiences such as the death of a comrade or having to handle dead body remains⁴⁷. Finally, less than half of Soldiers said that they would report a comrade’s unethical behavior⁴⁸. The reality is that humans are not perfect moral and rational soldiers and, in addition to empirical and practical mistakes, they frequently make moral ones too. In the fog of war, soldiers’ decision-making capabilities are often obstructed by feelings of revenge for lost comrades, weak leadership, dehumanization of the enemy, and poorly trained troops. Memoirs of drone pilots provide us with the perfect opportunity

⁴⁰ Dave Grossman, “On Killing: the Psychological Cost of Learning to Kill in War and Society,” (New York, NY: Open Road Integrated Media, 2014), 18.

⁴¹ Ibid., 18.

⁴² Ed Pilkington, “Life as a Drone Operator: ‘Ever Step on Ants and Never Give It Another Thought?’,” (The Guardian, Guardian News and Media, November 19, 2015, <https://www.theguardian.com/world/2015/nov/18/life-as-a-drone-pilot-creech-air-force-base-nevada>).

⁴³ See, Paul Scharre, interview with Lucas Perry, *AI Alignment Podcast: On Lethal Autonomous Weapons with Paul Scharre*, podcast audio, March 16, 2020, <https://podcasts.apple.com/us/podcast/the-future-of-life/id1170991978?i=1000468611509>.

⁴⁴ See, Office of the Surgeon Multinational Force-Iraq and Office of the Surgeon General (Army), “Final Report of the Mental Health Advisory Team (MHAT-IV),” (Washington, DC, 2006, <https://ntrl.ntis.gov/NTRL/dashboard/searchResults/titleDetail/PB2010103335.xhtml#>).

⁴⁵ Ibid., 35.

⁴⁶ Ibid., 36.

⁴⁷ Ibid., 38, 39, 40.

⁴⁸ Ibid., 37.

to compare the morality and rationality of soldiers when they are fighting remotely versus when they have skin-in-the-game. In the following chapter I ask whether RPAs make war more just and proceed to pinpoint the most appropriate role for the human component in warfare.

Chapter Three: RPAS - Overcoming Our Biological Inhibitions Against Killing

In 2005, a Predator missile struck its target near the border of Afghanistan. Drone pilot T. Mark McCurley was on duty at the Nellis Air Force Base. Conforming with the rules of engagement, the strike was an effort to prevent an imminent terrorist attack, and thus constituted a lawful act of self-defense. The Predator's target was a man known as the Facilitator, and intelligence gathered for over a month confirmed his status as legitimate target. He died during a routine phone call to his wife, the only pattern to which he regularly conformed. After identifying the man's scorched and dismantled body among the wreckage using the Predator's camera, McCurley drove home to cope alone⁴⁹. In a recent interview, McCurley described this incident as his most traumatic and intimate kill⁵⁰, and recounted the scene vividly, underscoring the moral injury that still impacted him years later. Drone pilots are physically distant from the scene, but they are emotionally hyper-connected in ways that oftentimes lead to long-term psychological trauma. The lack of skin-in-the-game can make remote pilots feel like they are "playing at God," and they often struggle to come to terms with the reality that they decide who lives and who dies without ever putting their own lives in danger. This chapter presents further examples of the greater intimacy of drone warfare, and highlights the moral injury suffered from not having skin-in-the-game. Then, I oppose the argument that says that skin-in-the-game is an inherent component of a just war, by showing that removing "fear of death" can reduce dehumanization and produce better moral agents. Finally, I demonstrate that the current remote model for war produces more ethical results and can be extended to remove humans from the decision-making process.

The largest part of a drone pilot's day is Intelligence, Surveillance, and Reconnaissance (ISR). A drone like the Predator has a flight-capacity of around 24 hours⁵¹ (22 hours longer than a regular piloted aircraft), and drone pilots will sometimes follow a specific target for days or even weeks at a time. During this extended spying, drone pilots become accustomed to their target's day-to-day life; one pilot even recalled witnessing a wedding⁵². In Bergen Rothenberg's collection of interviews, an anonymous drone pilot described the emotional investment characteristic of this kind of continued surveillance. He recounted: "you become immersed in their life. You feel like you are a part of what they're doing every single day. So, even if you're not emotionally engaged

⁴⁹ Mark T. McCurley and Kevin Maurer, "Hunter Killer: Inside America's Unmanned Air War" (New York, New York: Dutton, an imprint of Penguin Random House LLC, 2015), 130-136.

⁵⁰ See, "The Details of Drones, From A Pilot Who Flew Them," NPR (NPR, November 5, 2015), <https://www.npr.org/2015/11/05/454907560/the-details-of-drones-from-a-pilot-who-flew-them>.

⁵¹ Daniel Rothenberg, "Drone Wars: Transforming Conflict, Law, and Policy" (New York, NY: Cambridge University Press, 2015), 119.

⁵² See, "Drone", Video File, YouTube, posted by imacanism01, June 9, 2016, <https://www.youtube.com/watch?v=bMyPzoIMmKM>.

with those individuals, you become a little bit attached. [...] You see everything.”⁵³ Matt J. Martin described the vividness of the images and how he began to refer to his Predator drone as an extension of himself rather than a foreign object.⁵⁴ The emotional attachment that drone pilots develop towards what they are seeing is further reinforced by the images’ high definition and the heightened sense of nearness to the target⁵⁵.

Many people are quick to blame remote warfare for gamifying war and making it more trigger friendly, under the assumption that drone pilots are less emotionally aware and connected to the consequences of their actions. The claim that remote warfare can feel like a video game is not entirely unfounded, and in their memoirs, drone pilots often spoke of initially experiencing the images unfolding on their screens as surreal. When performing lethal actions, soldiers sometimes experience less moral conflict by using euphemisms, and distance from the conflict can have the initial effect of encouraging this type of behavior. With the Predator flying at 20,000 feet, the humanity of the moving dots on the screen can be difficult to recognize at first, and ex-serviceman Michael Haas denounced the use of colorful language such as “ants”⁵⁶ to describe indistinguishable targets, or “fun-sized terrorists” to indicate the children. In his memoir, drone pilot Matt Martin described his emotions during an air strike: “the man wasn’t *really* a human being. He was so far away and only a high-tech image on a computer screen. The moral aspect of it – that I was about to assassinate a fellow human being from ambush – didn’t factor in. Not at the moment. Not yet.”⁵⁷ In addition to this initial behavior, war, whether conventional or remote, involves deep emotional engagement; it is not uncommon for drone pilots to experience overwhelming feelings of vengeance and grief. The lag that exists between the moment a pilot calls a kill shot and the instant she is cleared to engage can mean the difference between saving innocent people or having to watch them die. When the latter outcome is realized, pilots often experience a host of debilitating emotions.

Unlike manned aircraft pilots who fly away after engaging their target, drone pilots are also responsible for providing damage assessment after a strike. Extremely high-quality images vivify the scenes of horror that pilots witness when searching through the wreckage in an attempt to identify the dismantled body of their target. In *Killer Robots*, Armin Krishnan comments on the evolving psychological impact of killing, noting that “apparently the zoom on the Foster-Miller SWORDS robot is so good that it allows the remote operators to even read the name tags of the soldiers they are going to kill over a distance of 300 feet between robot and target.”⁵⁸ McCurley notes the psychological impact of locking eyes with the Facilitator through the targeting pod seconds before he died. Because drone pilots also act as backup for soldiers on the ground, they sometimes have to identify the fallen figures of their friends, and in those instances the psychological strain on pilots can be even greater. Because of a drone’s ability to be everywhere,

⁵³ Daniel Rothenberg, “Drone Wars: Transforming Conflict, Law, and Policy” (New York, NY: Cambridge University Press, 2015), 115.

⁵⁴ Matt Martin and Charles W. Sasser, “Predator: The Remote-Control Air War over Iraq and Afghanistan: a Pilot’s Story,” (Minneapolis, MN: Zenith Press, 2010), 34.

⁵⁵ See, Press, Eyal. “The Wounds of the Drone Warrior.” The New York Times. The New York Times, June 13, 2018. <https://www.nytimes.com/2018/06/13/magazine/veterans-ptsd-drone-warrior-wounds.html>.

⁵⁶ See, Ed Pilkington, “Life as a Drone Operator: ‘Ever Step on Ants and Never Give It Another Thought?’,” The Guardian (Guardian News and Media, November 19, 2015), <https://www.theguardian.com/world/2015/nov/18/life-as-a-drone-pilot-creech-air-force-base-nevada>.

⁵⁷ Matt Martin and Charles W. Sasser, “Predator: The Remote-Control Air War over Iraq and Afghanistan: a Pilot’s Story,” (Minneapolis, MN: Zenith Press, 2010), 44.

⁵⁸ Armin Krishnan, “Killer Robots: Legality and Ethicality of Autonomous Weapons” (London, England: Routledge, 2016), 129.

there is no limit to how intimate the footage can be; for example, former drone pilot Brandon Bryant describes witnessing a particular target's burial ceremony.⁵⁹ Pilots witness more traumatic scenes of horror than the average soldier, and the vividness of an incident will sometimes haunt the killer years after it has happened. They often experience moral overload from the amount of lethal decisions that they have to make. In an interview, Brandon Bryant recalled the shock of discovering, upon retirement, that he had aided in the killing of 1, 626 people⁶⁰.

Compartmentalizing the visions that they see on the screen can be difficult among drone pilots, when the boundaries between mission and regular life are so hazy. Oftentimes, the transition between these two different worlds is extremely rapid, which makes it hard not to let the mission affect regular life, and vice versa. For example, one drone pilot recalled interrupting a visually jarring reconnaissance mission to take a lunch break with his wife, or to attend his child's soccer game⁶¹. Matt Martin describes his situation as living a "schizophrenic existence between two worlds, one as a combat pilot fighting a war halfway around the world, the other as an ordinary American citizen."⁶² Because remote fighting brings war closer to home than ever before, pilots lose the emotional support system that they would have had in traditional circumstances. Among soldiers deployed away from home, moral injury is countered partially by camaraderie and "unit cohesion"; this hermetic and structured environment provides soldiers with a space to cope with their lethal decisions. In addition, soldiers on the scene are sheltered from the societal backlash against warfare that accompanies any kind of conflict, which is not the case for drone pilots who have to interact with society in between their shifts. Wallace and Costello observe this phenomenon in their work, where they argue that "UAV pilots are exposed to unique stressors when they lose this unit cohesion and are instead 'deployed in place'."⁶³ On top of losing a sense of unit cohesion, there is not always a moral support system present in regular life for drone pilots to fall back on: most of a drone pilot's day-to-day is classified, and the guilt and shame surrounding many of their decisions can be psychological barriers to opening up. Drone pilots are in a constant state of limbo – not fully at war, and hardly at peace.

One way that soldiers justify lethal action is by appealing to their survival instinct. When soldiers have skin-in-the-game, it is easier to forgive lethal behavior because they are put in a situation where fighting is necessary in order to stay alive. But for drone pilots, whose lives are never at risk, this one-sided safety can feel cowardly and unjust. In an interview with NPR, Mark McCurley opened up about killing the Facilitator from the safety of his desk: "Even when I saw a guy shooting at troops and I helped take them out, for me, there was a moral justification that I was saving people's lives that were under attack, whereas this guy - this was kind of a him-and-me even though he never had a chance and it was never me."⁶⁴ The second chapter of this thesis

⁵⁹ Daniel Rothenberg, "Drone Wars: Transforming Conflict, Law, and Policy" (New York, NY: Cambridge University Press, 2015), 116.

⁶⁰ See, Ed Pilkington, "Life as a Drone Operator: 'Ever Step on Ants and Never Give It Another Thought?'," *The Guardian* (Guardian News and Media, November 19, 2015), <https://www.theguardian.com/world/2015/nov/18/life-as-a-drone-pilot-creech-air-force-base-nevada>).

⁶¹ Daniel Rothenberg, "Drone Wars: Transforming Conflict, Law, and Policy" (New York, NY: Cambridge University Press, 2015), 116.

⁶² Matt Martin and Charles W. Sasser, "Predator: The Remote-Control Air War over Iraq and Afghanistan: a Pilot's Story," (Minneapolis, MN: Zenith Press, 2010), 44.

⁶³ See, D. Wallace and J. Costello, "Eye in the sky: Understanding the mental health of unmanned aerial vehicle operators" [online], (*Journal of Military and Veterans Health*, Vol. 25, No. 3, Jul 2017: 36-41).

⁶⁴ See, "The Details of Drones, From A Pilot Who Flew Them," NPR (NPR, November 5, 2015), <https://www.npr.org/2015/11/05/454907560/the-details-of-drones-from-a-pilot-who-flew-them>.

concluded that dehumanization acts as a motivator to commit murder. Remote wars are fought very differently; drone pilots sometimes follow their targets for weeks on end and develop an attachment to their target unlike any experienced in traditional combat. In a recent New York Times Article⁶⁵, psychologist Shira Maguen reveals that veterans who were made to execute prisoners of war experienced higher rates of trauma than if they were executing strangers. The familiarity and understanding that pilots developed towards a target make it incredibly difficult to dehumanize them. Because skin-in-the-game and dehumanization are no longer justifiers for lethal actions, feelings of moral injury are precipitated, as pilots suffer from disillusionment and the injustice and futility of war. Michael Haas describes his day-to-day moral struggle thus: “you had to kill part of your conscience to keep doing your job every day - and ignore those voices telling you this wasn’t right.”⁶⁶

According to some, a responsibility gap arises with the use of AWS when blame is distributed across many people. Proponents of this theory argue that the individual feeling of moral responsibility is diluted, and it becomes easier to kill without feeling overwhelming guilt. This thesis is challenged by Paul Scharre in an NPR interview, where he talks about the unequal distribution of war trauma. According to him, drone pilots bear most of the burden for a kill because they emotionally engage with their target during the continuous and intimate surveillance period. The absence of skin-in-the-game produces a more *rational* and therefore *ethical* war, but also disillusiones pilots about the reality of war and causes them to suffer moral injury. Even though remote pilots do not suffer from physical brain trauma that comes with being in the battlefield, former pilots have affirmed suffering from many of the symptoms of PTSD. In their memoirs, pilots talk about losing sight of their moral compass, succumbing to depression, suffering from recurring nightmares and anxiety, and most of all feeling guilty and losing their sense of moral purpose. Killing is not a thoughtful and rational exercise, it is fundamentally unnatural and therefore requires enormous psychological and emotional sacrifices. With skin in the game and dehumanization, lethal action is not performed in a controlled and predictable manner because human beings succumb all too easily to fear and hatred. The increase in drone pilots’ ability to empathize with their targets makes them better moral agents than traditional soldiers, and remote fighting presents the opportunity to involve a plurality of opinions in a kill, which is more ethical than acting out of self-preservation and personal interest on the battlefield. However, drone pilots’ moral agency is wasted on this decentralized decision-making, because their individual opinion has less impact on the kill-decision. In order to salvage the psychological health of soldiers as well as maintain the level of ethicality observed in remote drone warfare, we must rethink the value of the human component in moral decision-making, and likely find a more suitable place for it.

There are many reasons why a human-in-the-loop architecture is not optimal for AWS, and, in *Army of None*, Paul Scharre mentions some of the pitfalls of only partially automating weapons systems like the Predator drone. One of the main complications is that AWS are weakened by the decentralized architecture – removing the human component from the physical battlefield disconnects the human “intelligence” from the AWS and exposes it to potential malicious attacks. If communications are jammed or compromised, the weapon is in danger of receiving foreign orders to fire on the wrong targets. In addition, the decision-making process in

⁶⁵ See, Eyal Press, “The Wounds of the Drone Warrior,” *The New York Times* (*The New York Times*, June 13, 2018), <https://www.nytimes.com/2018/06/13/magazine/veterans-ptsd-drone-warrior-wounds.html>).

⁶⁶ See, Ed Pilkington, “Life as a Drone Operator: 'Ever Step on Ants and Never Give It Another Thought?'," *The Guardian* (*Guardian News and Media*, November 19, 2015), <https://www.theguardian.com/world/2015/nov/18/life-as-a-drone-pilot-creech-air-force-base-nevada>).

remote fighting involves such a large array of opinions, that by the time drone pilots receive clearance for a shot, the opportunity has most likely passed. Machines' ability to follow a programmed set of rules can be an improvement on human soldiers when the programming implements the Laws of War, but there is no guaranty that a person with mal intentions would not implement non-ethical orders. An AWS with too little human control could result in the loss of human soldiers' ability to exercise free will and runs the risk of losing all military control to the wrong person. In order to avoid this scenario, there needs to be enough human involvement and diversity of opinion to ensure that one person, or a select group of individuals, are not all-powerful.

A human "on-the-loop" system provides a promising middle ground between keeping the human component in-the-loop and taking it fully out-of-the-loop. A human on-the-loop system is one in which the human has a supervisory role but is not directly involved in the decision-making process. This type of standard for AWS would provide a minimum of human control, and efficiently allocates and maximizes the strengths of both machines and humans. In order to minimize empirical, practical, and moral mistakes, the decision-making process should be performed by machines to allow for greater efficiency, less dehumanization, and avoid the moral injury endured by all types of combatants. In addition to the Laws of War, AWS that follow rules of engagement specific to the current mission, similarly to the ethical constrainer described in chapter 1, would allow for the abortion of any lethal action that disobeys this set of rules⁶⁷. The machine could also have a "dream" component with the ability to run mission simulations in a sandbox environment to better anticipate and prepare for unexpected situations and maximize the appropriateness of the lethal actions it generates in combat. Damage assessment, performed by humans, would form the basis for improving the system with human feedback. Continuing to have humans review the effects of war would serve to avoid the adiaphorization that might result from removing humans from the kill-decision.

Conclusion

The aim of this thesis is to disprove the anti-AWS argument that claims that having skin-in-the-game is a necessary component of a just war. In the preceding chapters, I contrast the killology of conventional warfare with that of remote drone fighting, commenting on the evolving psychological impact of killing. The logic-based programming and deep learning powering AI could enable a more just war by following the Laws of War with unprecedented precision. Thanks to advances in visual recognition, drones possess many of the necessary capabilities to autonomously navigate contested environments, perform advanced target recognition, and engage targets. The growing adoption of remotely piloted aircrafts in warfare has resulted in an impressive public backlash supporting many objections to the use of AWS. This thesis questions the validity of the anti-codifiability thesis, which states that moral rationality cannot be carried out by rule-based logic because morality is inexhaustive and requires contextualizing a moral dilemma. In Chapters two and three, I investigate the disconnect between the popular idealized version of war and soldiers' experiences. Interviews with soldiers deployed in Afghanistan reveal that when

⁶⁷ See, Ronald C. Arkin et al, "An Ethical Governor for Constraining Lethal Action in an Autonomous System:," (Fort Belvoir, VA: Defense Technical Information Center, January 1, 2009).

tensions are high, deployed soldiers tend to follow their instincts rather than comply with the Laws of War. Because humans have strong biological inhibitions against killing one another, violent conditioning and dehumanization are commonly found in conventional warfare settings. Skin-in-the-game can lead to unethical and unpredictable decision-making and does not permit as just of a war as society would like to believe. The memoirs of drone pilots allow us to comment on the psychological and ethical benefits of keeping pilots at a safe distance from the fight.

For the many reasons outlined in this thesis, removing the human component from the physical battlefield enables drone pilots to become better *moral agents* than if they had skin-in-the-game. My research points to the unique emotional connection that drone pilots develop towards their targets during the weeks of surveillance, allowing for a more profound immersion in a target's culture and daily habits. If moral decision-making requires Aristotelian rationality and involves contextualizing a moral dilemma, then drone pilots are better suited for the job than traditional soldiers. War is inherently irrational and reaches into some of the darkest corners of the human heart. We have seen that the process of dehumanizing acts as a kind of self-deception that helps soldiers overcome their biological inhibitions against lethal aggression. When soldiers realize the likeness of their foes, this dehumanization can culminate in feelings of moral injury; with remote fighting this disillusionment is simply accelerated. Memoirs of drone pilots reveal that removing the human component from the physical battlefield can eliminate this self-serving bias by enabling drone pilots to better understand their target and more easily adopt their point of view than if they were in an environment of propaganda, conditioning, and fear. This increase in each drone pilot's capacity to empathize with their target makes war more humane, but it also has debilitating consequences on pilots' mental health and can increase the number of *practical* and *empirical* mistakes that they make. Unique psychological stressors, induced by the loss of unit cohesion and camaraderie found in traditional battlefields, precipitate drone pilots' disillusionment with war. The remote position that was intended to ensure pilots' well-being produces feelings of cowardice when contrasted with their target's pathetic vulnerability.

The increase in cognitive combat intimacy observed in drone warfare is accompanied by a decentralization of the decision-making process of a given mission. The protocols in place involve a greater number of people than just the pilot herself, which ensures a better compliance with the Laws of War, and minimizes the moral mistakes made by frightened soldiers. However, this also means that drone pilots, while remaining incredibly affected by the outcome of a kill, do not have as great of an influence on their own actions as soldiers in the battlefield. The increase in drone pilots' moral agency is wasted in a decentralized decision-making process that maximizes the number of opinions involved in making a particular decision and is fixated on ensuring that each lethal action conforms with the Laws of War. Drone pilots experience moral injury as a result of having to perform actions that go against their moral judgment, and this moral injury can have negative outcomes on their performance. Having so many people participate in a kill increases the amount of time it takes to reach a decision, instead of capitalizing on machines' computational superiority. Furthermore, semi-autonomous weapons are fragile because their "intelligence" can come from thousands of miles away. An AWS with a human on-the-loop would take better advantage of machines' and humans' complementary strengths, helping to guarantee that cooler heads prevail, while minimize the amount of time taken for an AWS to make a normative decision, and act on it.

Bibliography

Books and Essays

- Arkin, Ronald C., Patrick Ulam, and Brittany Duncan. "An Ethical Governor for Constraining Lethal Action in an Autonomous System." Fort Belvoir, VA: Defense Technical Information Center, January 1, 2009.
- Banerjee S. 2018. A framework for designing compassionate and ethical artificial intelligence and artificial consciousness. *PeerJ Preprints* 6:e3502v2
<https://doi.org/10.7287/peerj.preprints.3502v2>.
- Bergen, Peter L., and Daniel. Rothenberg. *Drone Wars: Transforming Conflict, Law, and Policy* New York, NY: Cambridge University Press, 2015.
- Casey-Maslen, Stuart, Maziar Homayounnejad, Hilary Stauffer, and Nathalie Weizmann. *Drones and Other Unmanned Weapons Systems under International Law*. Leiden: Brill Nijhoff, 2018.
- Department of Defense Law of War Manual*. Washington, DC: General Counsel of the Department of Defense, 2015.
- Feldman, Philip, Aaron Dant, and Aaron Massey. "Integrating Artificial Intelligence into Weapon Systems." *ArXiv:1905.03899 [Cs]*, May 9, 2019.
<http://arxiv.org/abs/1905.03899>.
- Grossman, Lt. Col. Dave. *On Killing: The Psychological Cost of Learning to Kill in War and Society*. New York, NY: Open Road Integrated Media, 2014.
- Guarini, Marcello. "Introduction: Machine Ethics and the Ethics of Building Intelligent Machines." *Topoi* 32, no. 2 (October 2013): 213–15. <https://doi.org/10.1007/s11245-013-9183-x>.
- Jenkins, Ryan C., Michael Robillard, and Bradley Jay Strawser. *Who Should Die?: the Ethics of Killing in War* New York, NY: Oxford University Press, 2018.
- Keegan, John. *A History of Warfare*. New York: Alfred A. Knopf, 1997.
- Krishnan, Armin. *Killer Robots: Legality and Ethicality of Autonomous Weapons* London, England;: Routledge, 2016.
- Martin, Matt J., and Charles W. Sasser. *Predator: The Remote-Control Air War over Iraq and Afghanistan: a Pilot's Story* Minneapolis, MN: Zenith Press, 2010.

McCurley, T. Mark., and Kevin. Maurer. *Hunter Killer: Inside America's Unmanned Air War* New York, New York: Dutton, an imprint of Penguin Random House LLC, 2015.

O'Connell, Mary Ellen, Banning Autonomous Killing (August 21, 2013). Chapter in: *The American Way of Bombing: How Legal and Ethical Norms Change*, Matthew Evangelista & Henry Shue, eds. Ithaca, NY: Cornell University Press (2013, Forthcoming); Notre Dame Legal Studies Paper No. 1445. Available at SSRN: <https://ssrn.com/abstract=2313737>

Purves, Duncan, Ryan Jenkins, and Bradley Strawser. "Autonomous Machines, Moral Judgment, and Acting for the Right Reasons." *Ethical Theory and Moral Practice* 18, no. 4 (August 2015): 851–872.

Sassóli, Marco. "Autonomous Weapons and International Humanitarian Law: Advantages, Open Technical Questions and Legal Issues to Be Clarified" 90 (2014): 34.

Scharre, Paul. *Army of None: Autonomous Weapons and the Future of War* New York: W.W. Norton & Company, 2019.

Smith, David Livingstone. *Less Than Human: Why We Demean, Enslave, and Exterminate Others* 1st ed. New York: St. Martin's Press, 2011.

Wallace, D and Costello, J. Eye in the sky: Understanding the mental health of unmanned aerial vehicle operators [online]. *Journal of Military and Veterans Health*, Vol. 25, No. 3, Jul 2017: 36-41.

Wiegel, Vincent. "Wendell Wallach and Colin Allen: Moral Machines: Teaching Robots Right from Wrong: Oxford University Press, 2009, 273 Pp, ISBN: 978-0-19-537404-9." *Ethics and Information Technology* 12, no. 4 (December 2010): 359–61. <https://doi.org/10.1007/s10676-010-9239-1>.

Newspapers and Journals

"Know Thy Enemy: Better Understanding Foes Can Prevent Debilitating Hatred." Association of the United States Army, November 6, 2017. <https://www.USA.org/articles/know-thy-enemy>.

"Lethal Microdrones, Dystopian Futures, and the Autonomous Weapons Debate." *IEEE Spectrum: Technology, Engineering, and Science News*. Accessed May 27, 2020. <https://spectrum.ieee.org/automaton/robotics/military-robots/lethal-microdrones-dystopian-futures-and-the-autonomous-weapons-debate>.

- McCammon, Sarah. “The Warfare May Be Remote But The Trauma Is Real.” NPR. NPR, April 24, 2017. <https://www.npr.org/2017/04/24/525413427/for-drone-pilots-warfare-may-be-remote-but-the-trauma-is-real>.
- “Open Letter on Autonomous Weapons.” Future of Life Institute. Accessed April 29, 2020. <https://futureoflife.org/open-letter-autonomous-weapons/?cn-reloaded=1>.
- Pilkington, Ed. “Life as a Drone Operator: 'Ever Step on Ants and Never Give It Another Thought?'” The Guardian. Guardian News and Media, November 19, 2015. <https://www.theguardian.com/world/2015/nov/18/life-as-a-drone-pilot-creech-air-force-base-nevada>.
- Press, Eyal. “The Wounds of the Drone Warrior.” The New York Times. The New York Times, June 13, 2018. <https://www.nytimes.com/2018/06/13/magazine/veterans-ptsd-drone-warrior-wounds.html>.
- Scharre, Paul. “AI Alignment Podcast: On Lethal Autonomous Weapons with Paul Scharre”. *The Future of Life Institute*. Podcast audio. March 16, 2020. <https://podcasts.apple.com/us/podcast/the-future-of-life/id1170991978?i=1000468611509>.
- Smith, David Livingstone. “Less Than Human!: The Psychology Of Cruelty.” NPR. NPR, March 29, 2011. <https://www.npr.org/2011/03/29/134956180/criminals-see-their-victims-as-less-than-human>.
- “The Details of Drones, From A Pilot Who Flew Them.” NPR. NPR, November 5, 2015. <https://www.npr.org/2015/11/05/454907560/the-details-of-drones-from-a-pilot-who-flew-them>.
- “The Use of Armed Drones Must Comply with Laws.” ICRC, May 10, 2013. <https://www.icrc.org/en/doc/resources/documents/interview/2013/05-10-drone-weapons-ihl.htm>.
- Wallace, D and Costello, J. Eye in the sky: Understanding the mental health of unmanned aerial vehicle operators [online]. *Journal of Military and Veterans Health*, Vol. 25, No. 3, Jul 2017: 36-41.
- “We Should Not Ban ‘Killer Robots,’ and Here’s Why.” IEEE Spectrum: Technology, Engineering, and Science News. Accessed May 27, 2020. <https://spectrum.ieee.org/automaton/robotics/artificial-intelligence/we-should-not-ban-killer-robots>.
- “Why We Really Should Ban Autonomous Weapons: A Response.” IEEE Spectrum: Technology, Engineering, and Science News. Accessed May 27, 2020. <https://spectrum.ieee.org/automaton/robotics/artificial-intelligence/why-we-really-should-ban-autonomous-weapons>.

Official Reports

Office of the Surgeon Multinational Force-Iraq and Office of the Surgeon General (Army), Washington, DC. (2006). *Final Report of the Mental Health Advisory Team (MHAT-IV)*. <https://ntrl.ntis.gov/NTRL/dashboard/searchResults/titleDetail/PB2010103335.xhtml#>.

Slijper, Frank, and Pax (Organization). *Slippery Slope: The Arms Industry and Increasingly Autonomous Weapons*, 2019. <https://www.paxforpeace.nl/publications/all-publications/slippy-slope>.

Other Media

“Drone”. Video File. YouTube, posted by imacanism01, June 9, 2016, <https://www.youtube.com/watch?v=bMyPzoIMmKM>.

iai.co.il. Accessed May 3, 2020. <https://www.iai.co.il/p/harpy>.

“Nova.” Shield AI. Accessed May 4, 2020. <https://www.shield.ai/nova>.

"Slaughterbots". YouTube. 12 November 2017. Retrieved 29 April 2020.